Heuristic Algorithms for the Generalized Routing and Wavelength Assignment Problem

Esa Hyytiä*
Networking Laboratory
Helsinki University of Technology, Finland

Abstract

In this paper a generalized (static) routing and wavelength assignment problem is formulated, where the objective is to establish a given set of connections into the network while minimizing the number of wavelength channels at the same time. Connections can be normal lightpaths as well as anycast or multicast connections. Furthermore, each request can be for either a bidirectional or a unidirectional connection. In this context, we consider appropriately modified versions of the heuristic algorithms proposed in the literature. The performance of the algorithms is compared by means of numerical examples.

1 Introduction

The wavelength division multiplexing (WDM) technology has matured during the recent years and wavelength routed (WR) all-optical networks are currently been deployed into backbone networks. WR all-optical networks offer huge capacities, even terabits per second in a single fibre, and can thus meet the ever increasing capacity demands of Internet. In WR networks the nodes switch data optically per channel basis without any electro-optical conversions, which is cost effective and enables such huge speeds. At the same time each channel typically carries data at the speed of few Gb/s, which is suitable for electronic processing and routing.

Traditionally WR networks are seen as a layer, where a number of transparent lightpaths are established between the given s-d pairs. Establishment of a lightpath consists of choosing a route and a wavelength, which is not generally an easy task as even the wavelength assignment subproblem is NP-hard [1]. Each lightpath provides a logical link between its end nodes, i.e. at the logical layer lightpaths are used to transport data packets to their destinations.

The problem where a set of static lightpaths are to be established into a given network while minimizing the number of wavelength channels is often referred to as the static RWA problem [1, 2]. Alternatively one can consider dynamic traffic, where the lightpath requests obey some traffic pattern and the aim is to minimize, e.g. the blocking probability [3–5]. The static RWA problem can be extended to the logical topology design (LTD) problem by including the routing problem in the logical layer. In the LTD problem, the number of wavelength channels and the average traffic demands between each s-d pair are given, and the problem is to establish a set of lightpaths and determine the logical routes for the packet flows so that, e.g. the average packet delay is minimized [6–8]. Also the optical multicast trees have been considered in the

*Email: esa.hyytia@iki.fi

literature [1, 9]. Optical multicast tree can be used to broadcast data up to an entire channel from the root node to several destination nodes at the same time (unidirectional), or as a shared medium for several low bandwidth point-to-point flows by sharing the channel in time.

The third alternative is optical anycast, where a source node requests a lightpath to be established to any of the destination nodes [10]. Anycast requests can be seen as a generalization of normal lightpath requests. For example IPv6 networks provide anycast service. In WR network the optical anycast requests can be used, e.g. as an optimization tool to ensure survivability. The network designer can define anycast requests for certain nodes to ensure that there are alternative routes available to one of the main nodes. Alternatively, in operator to operator interface there can be a several border nodes through which all the traffic between the operators is to be delivered. In that case it is enough to find a free lightpath to one of the border nodes.

In this paper we assume that some intelligent party has come up with a set of requests consisting of normal lightpath, anycast (one-to-any) and multicast (one-to-many, i.e. a light-tree) requests and our aim is simply to set up the requests into the physical network with a minimal number of wavelength channels. We refer to this problem as a generalized RWA problem (GRWA) and propose heuristic algorithms to solve it by modifying existing algorithms proposed in the literature appropriately. We also propose layered algorithms for GRWA problem which set up the requested lightpaths and -trees on the layer by layer basis.

The rest of the paper is organized as follows. In Section 2 we formulate the GRWA problem. Then, in Section 3, a layered GRWA (L-GRWA) algorithm is presented. Then, in Section 4, several modified versions of previously proposed heuristic algorithms are presented. Section 5 contains numerical examples, and finally Section 6 finishes with the conclusions.

2 Problem Definition

We are given a network $G = (V, E)$, where each link consists of one or more unidirectional fibres. In particular, let $N = |V|$ denote the number of nodes and a $N \times N$ matrix $\mathbf{P}$ the physical topology, where p_{ij} represents the number of unidirectional fibres from i to j. Note that matrix $\mathbf{P}$ is symmetric when all links have the same number of fibres in both directions. Furthermore, we assume that the nodes cannot perform wavelength conversion and thus each lightpath/tree must use the same wavelength on each link along its path, known as *the wavelength continuity constraint*. Also, two connections sharing a same fibre must be on different wavelength channels, known as *the distinct channel assignment constraint* (DCA). For a given network one is suppose to establish a given set of connections consisting of 3 different types:

1. **unicast** lightpath requests (point-to-point)
2. **anycast** requests, i.e. the destination node can be any of the given set (one-to-any).
3. **multicast** requests, i.e. the optical signal is routed to several destinations (one-to-many).

Note that in multicast connection the set of links form a light-tree, while in the other cases the links form a lightpath. In each case the request can be for one or more channels, where each channel can be routed independently of others using a different route and wavelength. Furthermore, requests may be bidirectional, where the same route and wavelength is used in both directions, or unidirectional. Denote the set of connection requests with $\mathcal{A} = \{a_1, a_2, \ldots, a_n\}$. The generalized routing and wavelength assignment problem (GRWA) the objective is to establish the requests in $\mathcal{A}$ into a given network G with the minimum number of wavelength channels. without violating the wavelength continuity or the DCA constraints.

To this end let us first introduce some notation for request a:

$$\begin{aligned} t(a) &= \text{type of request (unicast, anycast or multicast)},\\ s(a) &= \text{source node of request } a,\\ D(a) &= \text{set of destination nodes of request } a,\\ m(a) &= \text{multiplicity, no. of wavelength channels},\\ b(a) &= \text{bidirectionality: } 2 \text{ (bidirectional) or } 1 \text{ (otherwise)}. \end{aligned}$$

3 Layered RWA Algorithms

In a network without wavelength conversion each lightpath must use the same wavelength on each link it travels through. In a static RWA problem, where one is asked to establish a set of lightpaths, the layered (graph) approach has turned out to work well (see [6, 11]). In layered approach one first considers the first wavelength layer λ_1 and assigns as many requests into it as possible. The order in which the feasibility of the path candidates is checked is crucial to the resulting configuration. It is easy to see, assuming that enough path candidates are available, that even the global optimum can be achieved with an appropriate order. Usually the longer paths (in number of hops) are harder to establish later than the shorter paths and thus RWA algorithms often try to allocate the longer paths first.

Previously presented layered RWA algorithms establish a given set of lightpaths into the network, but do not address the configuration of both anycast and multicast requests.[1] In this section we will describe two slightly different versions of the layered (graph) approach adapted to the GRWA problem, to which we refer to as L-GRWA and DL-GRWA algorithms. In rest of the paper we use the term “path” to denote a set of links which together fulfil a certain request and form a path or tree, respectively. The strategy of the layered algorithms is similar to the usual paradigm of RWA solutions, where at the first stage a set of possible routes for the requests is determined, and at the second stage the wavelength assignment is performed using a first-fit strategy (see Section 4 for two examples). The main distinction is that layered algorithms determine a route and wavelength at the same time instead of postponing the WA.

3.1 Layered GRWA Algorithm (L-GRWA)

In the first part L-GRWA algorithm determines one or more path candidates for each request and in the next step the requests are established into the network layer by layer using a subset of the predetermined paths. Generally an important factor in WR optical-networks is the number of links each request reserves. Thus, typically the physical length of the link can be neglected when determining the paths using a shortest path algorithm. To this end let us define two trivial mappings from the physical topology to “hop topology”. In particular, let $H_1(\mathbf{P}) = \mathbf{H}_1$ denote the unidirectional “hop matrix”, $H_2(\mathbf{P}) = \mathbf{H}_2$, respectively, the bidirectional “hop matrix”,

$$h_1(i,j) = \begin{cases} 1, & \text{if } p_{ij} > 0 \\ \infty, & \text{otherwise.} \end{cases} \qquad \text{and} \qquad h_2(i,j) = \begin{cases} 1, & \text{if } p_{ij} > 0 \text{ and } p_{ji} > 0 \\ \infty, & \text{otherwise.} \end{cases}$$

[1]Note that algorithms presented in [10] are solely for anycast requests.

For each request $a \in \mathcal{A}$ one or more "path" candidates are determined using the appropriate hop metric, $\mathbf{H}_1$ or $\mathbf{H}_2$, as follows:

1. For **unicast** (i.e. point-to-point) requests two parameters, (n_p, Δ_ℓ), are used to prune the path space, i.e. we include at most n_p shortest paths (in hops) such that no path may be more than Δ_ℓ hops longer than the shortest possible.
2. For **anycast** requests the algorithm determines, one or more, path(s) from the source node to each destination using the same path selection parameters as above.
3. For **multicast** requests (see, e.g. [1]) we determine a set of links by first generating the minimal spanning tree using Prim's algorithm, which is then iteratively pruned until no unnecessary leaf is left. (Alternatively one could use others heuristics, see e.g. [12].)

Thus, each candidate $x \in \mathcal{X}$ corresponds to exactly one request in $\mathcal{A}$ and specifies a set of links and the direction(s) the corresponding links are travelled. At this point it is not yet decided which of these paths will be actually used and on which wavelength channel.

At the second step the requested connections are established by executing a first-fit policy for one wavelength layer at a time. In particular, starting at the first wavelength layer, the algorithm sets up as many connections as possible using the paths in $\mathcal{X}$ at the current layer and then moves to the next layer. Each time a connection is set up the respective multiplicity $m(a)$ is decreased by one and once $m(a) = 0$ all the respective paths are removed from $\mathcal{X}$.

The order in which the candidates are tried is important for first-fit algorithms. To this end L-GRWA arranges the candidate list in a presumably favourable order. In particular, the "path" candidates are sorted using a three-level criterion defined by the vector $z(x)$,

$$z(x) = (\Delta_x,\ -\ell_x,\ -b_x),$$

where Δ_x is the number of extra hops path x uses when compared to the shortest path, ℓ_x is the number of used links and b_x is indicator of bidirectionality. The order of two paths is defined by the first unequal element in $z(x)$ (smaller first). The order ensures that at each step the algorithm sets up the longest possible route which uses the least amount of extra resources (i.e. extra hops). The algorithm is formally described in Algorithm 1.

3.2 Dynamic Layered GRWA Algorithm (DL-GRWA)

If we limit ourselves to the case of normal lightpath requests (i.e. a static RWA problem) the layered algorithm can be extended to adaptively use all possible paths. The term "dynamic" refers to the dynamic path selection at each step. Thus, instead of pruning the path space first the algorithm starts immediately to set up lightpaths into the first wavelength layer. Later it will be shown how anycast and multicast requests can be incorporated as well.

The dynamic layered GRWA (DL-GRWA) algorithm proceeds as follows. At each step DL-GRWA first picks one connection request and establishes a lightpath for it using the shortest free path at the current wavelength layer. Once no new connections can be established at the current layer the algorithm moves to the next layer and repeats the same procedure. This is repeated until all the requests have been set up. In order to reach the optimal configuration the algorithm should always pick "the right connection". We propose a heuristic order similar to the one used with L-GRWA, i.e. at each step the longest connection using the least amount of "extra links" is set up. This can be achieved as follows. Let $d(s, d)$ denote the length of the

Algorithm 1 Generalized Layered RWA, L-GRWA.

```
𝒳 ← path/tree candidates for each request
sort 𝒳 in increasing order of z(x) = ( Δ_x, −ℓ_x, −b_x)
W ← 0
while 𝒳 ≠ ∅ do
  W ← W + 1
  for each x ∈ 𝒳 do
    if links in x have free fibres at layer W then
      set up a lightpath/tree x at layer W
      decrease multiplicity m(a) by one
      if m(a) = 0 then
        remove x and other respective "paths" from X
      end if
    end if
  end for
end while
```

shortest path (in hops) from s to d in an empty network and, similarly, let $d'(s, d)$ denote the respective length at the current state. At each step the request $a = (s, d)$ to be established next is the one which minimizes the quantity,

$$c(s, d) = \left(\frac{N-1}{N}\right) \cdot d'(s, d) - d(s, d) = (d'(s, d) - d(s, d)) - \frac{1}{N} \cdot d'(s, d). \tag{1}$$

Clearly the first term, $d'(s, d) - d(s, d)$, corresponds to the number of extra hops and the second term, $\frac{1}{N} \cdot d'(s, d)$, is a normalized path length less than 1 (note that for all reachable (s, d)-pairs $d'(s, d) < N$). Thus, the number of extra hops, $d'(s, d) - d(s, d)$, serves as the primary key (less extra hops first) and the ties are broken by the length of the path (longer paths first).

The DL-GRWA algorithm is formally described in Algorithm 2 for a general case where $\mathcal{A}$ consists of both unidirectional and bidirectional requests. If there are both unidirectional and bidirectional requests one needs to determine the shortest paths (and the respective criteria) for both cases separately. Furthermore, an optional constraint on the maximum number of extra hops allowed, Δ_ℓ, has been introduced. Note that setting $\Delta_\ell = \infty$ means that all feasible paths are accepted. However, setting a finite limit on the number of extra hops might turn out to be useful when more connections are expected later.

Similarly, it is straightforward to extend the algorithm to handle anycast requests. For anycast request $a = \{s, D\}$ one must evaluate each possible destination and the path selection criterion must be adjusted slightly to take into account the alternative destinations. In particular, for anycast request a and $(s, d) \in a$ we suggest using the shortest path in an empty network to the nearest node from D as the reference length,

$$c(a, s, d) = \left(\frac{N-1}{N}\right) \cdot d'(s, d) - \min_{i \in D(a)} d(s, i). \tag{2}$$

Multicast connections could be handled with dynamical routing algorithm as well, but that would increase the complexity considerably. Thus, we propose using a fixed set of routes for multicast connections along the lines of L-GRWA algorithm and set up them first before

Algorithm 2 Dynamic Layered RWA, DL-RWA.

```
run all-pairs shortest path algorithm for H_1(P) and H_2(P) ⇒ distances {D_1, D_2}
W ← 1
P' ← P
while A ≠ ∅ do
  run all-pairs shortest path algorithm for H_1(P') and H_2(P') ⇒ distances {D'_1, D'_2}
  C_i ← ((N−1)/N) · D'_i − D_i, i = 1, 2
  set c(a) = (C_{b(a)})_{s(a),d(a)}  ∀ a ∈ A {path selection criteria}
  A' ← {a ∈ A : c(a) < Δ_ℓ − 1}
  if A' = ∅ then
    W ← W + 1
    P' ← P
  else
    a ← arg min_{a∈A'} c(a)
    set up request a using the respective shortest path p to the current layer W
    reduce the number of free fibres in P' along path p
    decrement multiplicity m(a) by one
    if m(a) = 0 then
      remove request a from A
    end if
  end if
end while
```

continuing with DL-GRWA. It is assumed that this approach does not hinder the final solution much as long as the proportion of multicast requests is reasonably small. It is easy to show that the complexity of DL-GRWA is $|V|^3 \cdot |\mathcal{A}|$.

4 Two Stage Algorithms

For comparison purposes we describe next static RWA algorithms proposed in the literature. All these algorithms belong to a class of "two stage algorithms" (TSA), where the name emphasizes the fact that they solve the RWA problem in two stages. In particular, in the first stage the algorithms determine a path for each request, and in the second stage each path is assigned a feasible wavelength channel. Note the distinction to the (D)L-GRWA algorithms, which select the path at the same time with the wavelength channel.

All these TSA algorithms are originally designed to solve a special case like a single fibre network with bidirectional unicast requests. In order to overcome such limitations the algorithms have been modified appropriately. In particular, anycast and multicast requests need special treatment. Each algorithm has been extended to be able to continue from a given (fixed) configuration, i.e. to set up the remaining requests into the network while keeping the currently active lightpaths and -trees fixed. Thus, if a given algorithm is not capable of handling multicast requests, then such requests are first established into the network using L-GRWA, similarly as with DL-GRWA. In cases where an algorithm is unable to handle anycast requests such requests are first converted to normal unicast requests by choosing, in each case, the nearest node from the set of possible destinations.

Algorithm 3 Wavelength Assignment Step [10, 13]

for each $x \in \mathcal{X}$ in the decreasing order of the number of hops **do**
  set up x using the smallest feasible λ-channel
end for

Furthermore, possible multifibre cases must be taken care of. The path selection process could often ignore the presence of multiple fibres, especially if the number of fibres is constant on all links. Wavelength assignment step, however, should be aware of multiple fibres and act appropriately. In the used WA algorithm the paths are first sorted in the decreasing order of their lengths (in hops). Then, in that order, each path is assigned the lowest feasible wavelength channel, i.e. so called first-fit policy. The wavelength assignment step is formally described in Algorithm 3. Note that the assignment can be performed one wavelength layer at a time basis as well, which makes the implementation somewhat easier.

4.1 Shortest Path Approach (SP)

A straightforward approach for determining the paths is to use the shortest path (SP) in terms of optical hops for each request and then assign the wavelength channels using, e.g. WA Algorithm 3. This method gives usually reasonably good solutions when physical and logical topologies are similar, but can fail badly if there is a substantial mismatch between them. In this paper SP algorithm is used to set up unicast requests, while possible any- and multicast requests are handled as described above. Furthermore, possible multiple fibres are ignored in the routing phase, i.e. algorithm does not favour routes with multiple fibres on some link(s).

4.2 Minimum-Hop Heuristics (MNH)

The next TSA algorithm is minimum-hop heuristic (MNH) proposed by Baroni et al. in [13]. MNH is a greedy algorithm which first tries to lower the maximum congestion in the network before proceeding with the WA step. The initial set of paths is equivalent to the paths of the SP approach described above. After that the algorithm tries to find an alternative path for one request at a time leading to a lower maximum congestion along the particular path.

The MNH algorithm is designed to establish only normal lightpath requests and thus cannot handle anycast or multicast requests. As mentioned earlier, these limitations are overcome by configuring the possible multicast requests first using the L-GRWA algorithm and transforming possible anycast requests to normal lightpath requests to the nearest anycast destination (in hops) before proceeding with MNH algorithm.

Furthermore, the original MNH algorithm is generalized in three ways. Firstly, the minimum hop constraint has been relaxed in the outer loop by using an additional parameter Δ_ℓ, which defines the maximum number of extra hops allowed for paths. Secondly, the algorithm has been extended to multifibre case by defining the link loads to be the smallest integer not less than the number of users u_{ij} divided by the number of fibres p_{ij},

$$C_{ij} = \lceil u_{ij}/p_{ij} \rceil.$$

Thirdly, the algorithm takes into account already set up connectiongs $\mathcal{X}_{\text{fixed}}$.

Algorithm 4 Relaxed Minimum-Hop Heuristic, MNH+ [13].

```
for each (unicast) request a ∈ 𝒜 do
  find (any) minimum-hop path for a and store (m(a) copies of it) to 𝒳
end for
compute the initial link loads C_ij = ⌈u_ij/p_ij⌉ based on the paths in 𝒳 (and possibly 𝒳_fixed)
for K = 0, ..., Δ_ℓ do
  repeat
    for each route x ∈ 𝒳 do
      find the shortest path p which reduces congestion on the most loaded link
      if length(p) is at most K hops longer than the minimum-hop path then
        substitute the current path by the alternative path and update link loads
      end if
    end for
  until no further substitutions are possible
end for
assign wavelengths to routes in 𝒳 using Algorithm 3
```

The relaxed MNH algorithm (MNH+) is formally described in Algorithm 4. Note that the original version proposed in [13] is equivalent to MNH+ with $\Delta_\ell = 0$. We have found out that allowing longer paths tends to give considerably better results than the original version. However, it is also possible that in some instances relaxing the minimum-hop constraint leads to a harder WA problem and consequently more wavelength channels are required in the end.

4.3 TJWW Algorithms

The last algorithm in this family is an anycast RWA algorithm proposed in [10], which we refer to as TJWW algorithm. The algorithm first determines a set of routes and the next stage is the WA step using a standard greedy approach[2] described in Algorithm 3. As TJWW cannot handle multicast requests we handle such first using L-GRWA algorithm before proceeding with TJWW. Unicast requests are treated as trivial anycast requests.

The TJWW algorithm has two routing parameters, α and β. The first parameter α is used to adjust path selection probabilities per destination node based on the number of hops it takes to reach the respective destination. First the shortest paths are determined to each possible destination and then the destination is chosen randomly using the weights $w(x) = R \cdot {\ell_x}^{-\alpha}$, where ℓ_x is the number of hops in x and R a normalization constant. Note that when $\alpha = 0$ each destination is chosen with an equal probability and when $\alpha \to \infty$ (one of) the closest destination(s) is chosen. The second parameter β determines whether the maximal link load is controlled during the path selection process ($\beta = 1$) or not ($\beta = 0$). If the maximum link load is controlled, one first determines a "residual" graph G' from G by removing links with $C_{\max}$ or more connections. Then the shortest paths are determined using G'. If no path is found then $C_{\max}$ is increased by one. Initially $C_{\max}$ is set to 1. Hence, the final value of $C_{\max}$ serves as a lower bound for the number of wavelength channels needed.

Also the TJWW algorithm must be modified for our purposes. Firstly, the extension to multifibre case can be made in several ways, from which we have chosen perhaps the most

[2] WA step is not clearly described in [10] but it is assumed to be similar to this. Also it was assumed that $m(a) = 1$.

Algorithm 5 Two Stage Anycast Algorithm TJWW proposed in [10].

set $C_{\max} = 1$ and compute initial link loads C_{ij} based on $\mathcal{X}_{\text{fixed}}$
while $\mathcal{A} \neq \emptyset$ **do**
  select $a \in \mathcal{A}$ and reduce its multiplicity $m(a)$ by one
  if $m(a) = 0$ **then**
    remove a from $\mathcal{A}$
  end if
  repeat
    set $\mathbf{H} \leftarrow H_{b(a)}(\mathbf{P})$
    if $\beta = 1$ **then**
      remove all links (i, j) with $C_{ij} \geq C_{\max}$ in $\mathbf{H}$
    end if
    $\mathcal{P} \leftarrow$ shortest path in $\mathbf{H}$ from $s(a)$ to each $d \in D(a)$
    if $\mathcal{P} = \emptyset$ **then**
      $C_{\max} = C_{\max} + 1$
    else
      pick path $x \in \mathcal{P}$ using weights $w(x) = R \cdot {\ell_x}^{-\alpha}$, where R is a normalization constant
    end if
  until a path is found
  add path x into path set $\mathcal{X}$
  set $C_{ij} \leftarrow C_{ij} + 1/F_{ij}$ for all $(i, j) \in x$
end while
assign wavelengths using Algorithm 3

intrinsic one, i.e. the link loads are defined as the number of users divided by the number of links. The maximum link load $C_{\max}$ is still increased in steps of one. Secondly, the directionality must be explicitly taken care of when determining the shortest paths. Thirdly, the fixed connection must be taken into account by adjusting the initial link loads appropriately. The procedure is formally described in Algorithm 5. In [10] four parameter sets we used from which so called BWC algorithm, where $\alpha = \beta = 1$, was reported to give the best results. Thus, in Section 5 we will focus on BWC and compare its performance to other algorithms.

5 Numerical Results

In this section we present some numerical results. The path selection parameters used in L-GRWA are $n_p = 4$ and $\Delta_\ell = 1$, i.e. at most 4 paths are included and each path may not be more than 1 hop longer than the shortest possible. From TJWW algorithms we consider BWC (i.e. $\alpha = \beta = 1$). Hence, the maximum link load is controlled and anycast destinations reachable with fewer hops are preferred in the path selection process. From MNH algorithms we consider both basic MNH ($n_p = \infty$ and $\Delta_\ell = 0$) and MNH+ with $\Delta_\ell = 2$, i.e. the alternative paths are allowed to consist of at most zero or 2 extra hops, respectively.

5.1 MCI-ISP network

The first set of simulations uses MCI ISP backbone network depicted in Fig. 1 consisting of 19 nodes and 32 links. All links are assumed to have a single fibre in both directions and no

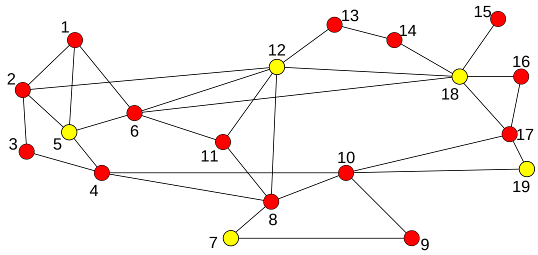

Figure 1: Example network: MCI backbone network in United States [10].

wavelength conversion is available. In order to see the relative performance of the algorithms we consider four different traffic scenarios each consisting of total 4 million requests.

1. **Unicast Requests**: a set of n unicast (i.e. lightpath) requests between random node pairs.
2. **Anycast Requests**: The anycast scenario has been adopted from [10]. In this scenario nodes 5, 7, 12, 18 and 19 serve as a single anycast destination set D and (bidirectional) anycast requests originate from the rest of the nodes randomly, i.e. each request is a bidirectional anycast request from some node $s \notin D$ to any node in D.
3. **Heterogeneous Requests**: The anycast scenario described above is not very realistic. In order to get a better idea how different algorithms perform in more realistic situations we consider also a scenario with a mixture of normal and anycast requests where both the source and the destination node are randomly picked, $s \in V \setminus D$ and $d \in V \setminus \{s\}$. If the destination node d belongs to set D, then the request is interpreted as an anycast request to set D, otherwise as a normal lightpath request to node d.
4. **Heterogeneous Requests Including Multicast**: The final example is the most demanding one, i.e. the set of requests includes the normal, anycast and multicast requests. The request sets were generated as follows. For each request first the source node $s \in V$ is picked randomly. If $s \in D$ then the request is interpreted as a bidirectional multicast request from s to $V \setminus D$. Otherwise we pick a random destination node $d \in V \setminus \{s\}$ and if $d \in D$ then the request is interpreted as a bidirectional anycast requst from s to D and otherwise as a normal lighpath request from s to d.

For each traffic scenario we proceed as follows. At start the network is empty and a set of n requests is generated. Then the chosen algorithm sets up the requests and the number of wavelength channels used is recorded. The average number of used wavelength channels is estimated by running 10000 independent request realizations. The numerical results can be seen in Table 1. In all scenarios the SP algorithm with first-fit WA has the poorest performance, as expected. Furthermore, it can be seen that the layered algorithms seem to work very well and that the DL-GRWA algorithm performs slightly better than L-GRWA algorithm. The improvement with DL-GRWA algorithm when compared to the best result obtained by SP, BWC or MNH(+) in each case is order of $2-4\%$ for unicast requests, $11-15\%$ for anycast requests, $0.5-2\%$ for combination of uni/any-cast requests and $2-3\%$ in case of all kinds of requests. Thus, in anycast case the DL-GRWA seems to be superior to other algorithms.

Also MNH+ algorithm performs well except in the anycast scenario where BWC and (D)L-GRWA algorithms obtain better results on average. This is probably due to the adaptive se-

request set	nr. of requests	SP	BWC	MNH	MNH+	L-GRWA	DL-GRWA
unicast	20	5.17	4.16	4.50	3.89	4.15	3.75
	40	8.72	7.16	7.47	6.66	6.98	6.49
	60	12.09	9.97	10.28	9.31	9.67	9.10
	80	15.40	12.67	12.98	11.88	12.29	11.66
	100	18.68	15.31	15.62	14.41	14.88	14.19
anycast	20	4.09	2.84	3.77	2.91	2.42	2.42
	40	6.86	4.79	6.14	4.87	4.11	4.10
	60	9.57	6.58	8.34	6.75	5.75	5.74
	80	12.27	8.32	10.48	8.58	7.36	7.34
	100	14.98	10.02	12.56	10.39	8.94	8.90
uni/any-cast	20	4.71	3.99	4.14	3.62	3.89	3.56
	40	7.93	6.86	6.90	6.23	6.63	6.17
	60	11.05	9.55	9.50	8.79	9.28	8.72
	80	14.10	12.16	12.03	11.30	11.83	11.24
	100	17.15	14.72	14.54	13.80	14.39	13.75
uni/any/multi-cast	20	8.64	8.16	8.10	7.82	7.94	7.55
	40	16.19	15.43	15.16	15.03	15.06	14.55
	60	23.71	22.56	22.19	22.19	22.13	21.55
	80	31.24	29.63	29.22	29.35	29.19	28.57
	100	38.75	36.67	36.24	36.49	36.24	35.60

Table 1: Numerical results with MCI-ISP network.

lection of the anycast destination in BWC and (D)L-GRWA algorithms. Note that most of the connections are rather short due to the fact that anycast destinations are quite dense.

The heterogeneous case consisting of a mixture of uni/any-cast requests gives similar results as the unicast case, but the difference between the algorithms is smaller, which suggests that the BWC algorithm handles anycast requests less efficiently than unicast requests, perhaps because the destination node is chosen impractically in some cases. Furthermore, the L-GRWA algorithm seems to suffer from the finite number of routes available (with the used parameters). Only when the number of requests increases the difference to BWC algorithm becomes larger.

Note that only L-GRWA algorithm sets up multicast requests as a part of its normal execution, while the other algorithms rely on L-GRWA first and let it set up the multicast requests for them. Indeed, in the last scenario including also multicast requests it can be noted that the (D)L-GRWA and MNH(+) algorithms have almost identical performance. Especially interesting is that in last three cases MNH outperforms MNH+, i.e. even the maximum link load gets lower the WA step has become harder and more wavelength channels are used. The fact that differences are small in the last set is probably partly due to the fact that about 26% of the requests are multicast requests and they are handled by the same algorithm. The small improvement DL-GRWA was able to find when compared to others must have happened at the later stages of the algorithm when the remaining lightpaths were set up.

5.2 NSFNET and EON

In [14–16] the authors have studied a slightly different problem, i.e. a problem where the number of available wavelength channels is given and the goal is to establish as many lightpaths of a given virtual topology (VT) as possible (or to be exact, to carry as much point-to-point traffic as possible). As the number of available wavelength channels is increased a point is reached where a given algorithm manages to establish all the requested lightpaths. This is equivalent to the problem studied in this paper, i.e. what is the minimum number of wavelength channels needed to establish a given set of lightpaths into the network.

network	case	SP	BWC	MNH	MNH+	L-GRWA	DL-GRWA	RWABOU [16]	KS [15]
NSFNET,	uni.	25	20	21	19	22	20	20	23
NSFNET,	bi.	32	36	28	26	30	28	$\star$	NA
EON,	uni.	46	20	20	19	21	18	22	$\star$
EON,	bi.	53	27	25	25	26	25	$\star$	NA
EON II,	uni.	28	17	18	16	18	16	NA	NA
EON II,	bi.	35	23	24	23	23	23	NA	NA

Table 2: The number of wavelength channels required to establish virtual topologies given in [14, 16] using either unidirectional or bidirectional lightpaths.

Krishnaswamy et al. in [14,15] presented a formulation based on LP-relaxation and considered only unidirectional cases. The resulting algorithms (KS) are considerably more complex than TSA or layered algorithms. Jaumard et al. in [16], on the other hand, considered both unidirectional and bidirectional cases. They proposed an algorithm called RWABOU, which is a two phase algorithm where the routing and the wavelength assignment phases are repeated until no further progress is obtained. Thus, one round in RWABOU is, in some sense, equivalent to whole TSA. The possible routes are chosen in a similar fashion than with L-GRWA, i.e. n_p shortest paths for each (s, d)-pair are first determined. Then, using different rerouting strategies together with a wavelength assignment phase based on Tabu search the final solution is obtained. Thus, also this algorithm is considerably more complex than (D)L-GRWA and TSA algorithms described in Sections 3 and 4. Following [14–16], we consider two example networks, EON and NSFNET, depicted in Fig. 2 The three test cases are:

- NSFNET: 14 nodes and 21 bidirectional single-fibre links
- EON: 20 nodes and 39 bidirectional single-fibre links
- EON II: otherwise the same as EON, but links Madrid-Paris, Paris-Milan, Milan-Berlin, Berlin-London and London-Paris are upgraded to consist of two bidirectional fibres

The bottleneck in EON tends to be the two links connecting Lisbon and Madrid to the rest of the network. Thus, Madrid-Paris and some other links have been upgraded in EON II, which is illustrated by thicker lines in Fig. 2. Note that NSFNET and EON are the same networks as in [14–16], from which also the used traffic matrixes defining the sets of lightpaths originate from (268 unidirectional connections in case of NSFNET and 374 in case of EON/EON II). Also the bidirectional case is obtained the same way as in [16], i.e. the number of bidirectional channels requested between s and d is

$$\mathrm{T}_{\mathrm{bi}}(s, d) = \max\{\, \mathrm{T}_{\mathrm{uni}}(s, d),\ \mathrm{T}_{\mathrm{uni}}(d, s)\,\}.$$

The numerical results with the algorithms discussed in this paper as well as the results obtained in [14–16] are presented in Table 2. Note that RWABOU and KS algorithms are not able to deal with a multifibre case, which is indicated by NA in the respective columns. Also it is worth noting that Jaumard et al. in [16] increased the number of available wavelengths in steps of 2, so the results corresponding to RWABOU could be one wavelength less. Furthermore, they considered only up to 24 wavelength channels and thus a symbol $\star$ in RWABOU column indicates that the respective solution requires more than 24 wavelengths. Krishnaswamy et al. [14, 15] studied only unidirectional case and up to 25 wavelength channels, i.e. $\star$ in KS column indicates that more than 25 wavelength channels are required.

From the results it can be seen that both MNH(+) and (D)L-GRWA algorithms perform surprisingly well. This is probably due to the multiple fibres in some links, which are equivalent

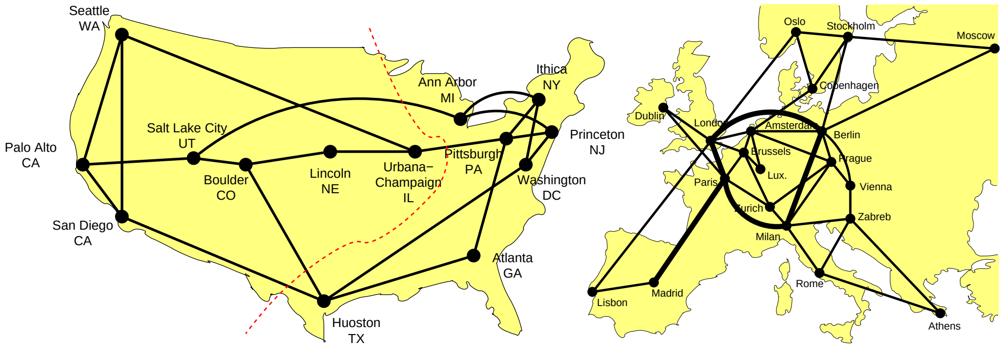

Figure 2: Example networks NSFNET and EON [2, 13, 14, 16].

to partial wavelength conversion. Even though RWABOU and KS algorithms are considerably more complex they came up with less efficient configurations in all cases. Furthermore, without appropriate modifications they are limited to unicast connections in a single fibre case. Note that MNH(+) algorithm iterates the routing based on link loads in a greedy manner and then performs the WA step using a fast heuristic algorithm, while the (D)L-GRWA algorithms come up with a solution in a single round. Also it can be seen from the results that EON II seems to be easier problem than EON, e.g. the simple SP algorithm manages to find a considerably better configuration for EON II than for EON.

6 Conclusions

In this paper we have formulated a generalized static RWA problem, where the objective is to establish a given set of uni/any/multi-cast requests into a WR optical network using a minimum number of wavelength channels. For solving the GRWA problem several heuristic algorithms were considered. First two versions of a layered heuristic algorithm using a first-fit strategy in a novel order for one wavelength layer at a time were proposed. Then, several other previously proposed algorithms were extended appropriately. The performance of the algorithms was evaluated by numerical examples. In all test cases of the first test set with MCI-ISP network the DL-GRWA algorithm achieved the lowest mean number of wavelength channels. In the special case consisting of only anycast requests the layered algorithms were shown to be superior to the previously proposed anycast TJWW algorithms [10]. In the next set of tests, adopted from the literature, the problem was to establish a given virtual topologies consisting of numerous lightpaths. In this case both MNH+ and DL-GRWA algorithms turned out to be superior to the other algorithms, including considerably more complex algorithms RWABOU and KS. The presented heuristic algorithms are relatively fast and robust and thus can be used in a "what if" analysis where different options are evaluated at the higher layer. Furthermore, the L-GRWA approach can be used when network survivability is considered by choosing the alternative routes for anycast destinations so that they do not coincide with the "primary" routes and thus ensure the survivability in case of a fibre cut.

References

[1] George Rouskas. Routing and wavelength assignment in optical WDM networks. In John Proakis, editor, *Wiley Encyclopedia of Telecommunications*. John Wiley & Sons, 2001.

[2] Stefano Baroni. *Routing and wavelength allocation in WDM optical networks*. PhD thesis, University College London, May 1998.

[3] S. Subramaniam and R.A. Barry. Wavelength assignment in fixed routing WDM networks. *IEEE International Conference on Communications 1997*, pp. 406–410, 1997.

[4] A. Mokhtar and M. Azizoğlu. Adaptive wavelength routing in all-optical networks. *IEEE/ACM Trans. Networking*, 6(2):197–206, Apr. 1998.

[5] E. Hyytiä and J. Virtamo. Adaptive importance sampling in routing and wavelength assignment. *European Trans. Telecommunications (ETT)*, 13(4):331–339, Jul./Aug. 2002.

[6] Zhensheng Zhang and Anthony S. Acampora. A heuristic wavelength assignment algorithm for multihop WDM networks with wavelength routing and wavelength re-use. *IEEE/ACM Trans. Networking*, 3(3):281–288, Jun. 1995.

[7] Rudra Dutta and George N. Rouskas. A survey of virtual topology design algorithms for wavelength routed optical networks. Technical Report TR-99-06, Department of Computer Science, North Carolina State University, 12, 1999.

[8] Nicolas Puech, Josué Kuri, and Maurice Gagnaire. Models for the logical topology design problem. In *Networking 2002*, pp. 301–313, Pisa, Italy, May 2002.

[9] D.-N. Yang and W. Liao. Design of light-tree based logical topologies for multicast streams in wavelength routed optical networks. In *Infocom '03*, pp. 32–41, Mar. 2003.

[10] Maoning Tang, Weijia Jia, Hanxing Wang, and Jianxin Wang. Routing and wavelength assignment for anycast in WDM networks. In *International Conference on Wireless and Optical Communications (WOC)*, pp. 301–306, Banff, Alberta, Canada, Jul. 2003.

[11] C. Chen and S. Banarjee. A new model for optimal routing and wavelength assignment in wavelength division multiplexed optical networks. In *Infocom'96*, pp. 164–171. 1996.

[12] S. Ramanathan. Multicast tree generation in networks with asymmetric links. *IEEE/ACM Trans. Networking*, 4(4):558–568, Aug. 1996.

[13] S. Baroni and P. Bayvel. Wavelength requirements in arbitrarily connected wavelength-routed optical networks. *IEEE J. Lightwave Technol.*, 15(2):242–251, Feb. 1997.

[14] R.M. Krishnaswamy. *Algorithms for Routing, Wavelength Assignment and Topology Design in Optical Networks*. PhD thesis, Indian Institute of Science, Bangalore, India, 1998.

[15] R.M. Krishnaswamy and K.N. Sivarajan. Algorithms for routing and wavelength assignment based on solutions of LP-relaxations. *IEEE Communications Letters*, 5(10):435–437, Oct. 2001.

[16] B. Jaumard and T. D. Hemazro. Routing and wavelength assignment in single hop all optical networks with minimum blocking. Research Report G-2004-12, Université de Montréal, Groupe d'Etude et de Recherche en Analyse des Décisions, Feb. 2004.